SANDVINE

Cloud Evolved ANI

Unlocking the promise of dynamic, virtualized deployments

73%

of telecommunications financial executives expect emerging new technologies, devices, and services to have a positive impact on their business over the next year.

Source: KPMG, "2014 Media & Telecommunications Industry Outlook Survey," December 2014.

Sandvine, Red Hat, Intel, Openet, Amartus, and Cobham present a readily deployable, reliable, robust, BSS-enabled NFV solution, allowing for virtual network infrastructure monetization and business model transformation for the future

INTRODUCTION

The dynamic capabilities of cloud-based, virtualized networks hold much promise for network operators:

- Improved service agility for introducing new services sooner and to experience faster time-to-value with those services
- Reduced cost of launching new services, allowing operators to bring to market experimental services that were previously infeasible
- New deployment possibilities, particularly for mobile and satellite operators
- Operational efficiencies and lighter initial investments will increase cost-effectiveness

At the same time, network users will benefit from higher quality of experience (QoE) and a wider range of valued services enabled by advanced technologies like Sandvine Application and Network Intelligence (ANI).

Looking beyond applications to enhance today's networks, it is well-known that 5G architecture is cloud-centric. Hence, cloud infrastructure support is a pre-requisite for any independent software vendors to offer solutions for 5G networks—specifically the 5G Multi-access Edge Computing (MEC) and 5G Stand-Alone (SA) networks.

Likewise, complete support for deployments on Infrastructure-as-a-Service (IaaS) environments like Amazon Web Services (AWS) is a key platform requirement for any vendors aspiring to become part of the 5G future.

When examining cloud solutions, it is helpful to think of different architectural layers:

- Application Layer: Applications packaged as Virtual Network Functions (VNFs) or Cloud-native Network Functions (CNFs)
- Management Layer: Virtualized Infrastructure Management (VIM), with OpenStack being the most common option for private cloud implementations
- Physical Layer: Physical Infrastructure or Network Function Virtualization Infrastructure (NFVI)

In practice, hyperscale providers including AWS, Microsoft Azure, and Google Cloud Platform (GCP) supply both VIM and NFVI.

5G CLOUD ARCHITECTURE REQUIREMENTS

The promise of 5G is tremendous, but meeting the expected performance improvements, achieving magnitude of scale required, and creating new services to leverage these capabilities will be difficult to fulfill.

Innovations including network functions virtualization (NFV), software-defined networking (SDN), MEC, and network slicing individually and collectively provide part of the answer, but they are themselves dependent on specific architectural requirements, including:

- Control and User Plane Separation (CUPS)
- Scalable stateless architecture
- Service-based interfaces

When implemented correctly, hyperscale clouds offer the ideal infrastructure for 5G networks.

Control and User Plane Separation

Under the 3GPP standard, the control plane and the user plane are separated. Within this architectural paradigm, the User Plane Function (UPF) handles packet processing, while signaling and other functionality is handled in the control plane by the Session Management Function (SMF).

The typical architecture has the control plane in the core network, with data plane functions located at (or close to) the edge.

Most hyperscale cloud providers have implemented a unified core-to-edge hierarchy, which connects all services for a single tenant (i.e., an operator).

Scalable Stateless Architecture

A stateless architecture helps to create reliable, scalable, and resilient networks. Hyperscale providers typically decouple the application layer from the storage layer, thereby contributing to a stateless and scalable design.
Within such a deployment, operators are able to select from a range of data storage options such as NoSQL databases, columnar databases, relational databases, key-value and object storage, and in-memory caches such as Redis and Memcached.

Service-Based Interfaces

In a 5G network, arguably the most significant difference with previous generations is the ability to provision service-based interfaces using HTTP/2.
Having these interfaces not only helps with inter-service communications, but greatly extends interoperability with native services offered by the hyperscale providers such as API gateways, service meshes, queuing and notification services, and container orchestration APIs for Kubernetes and container-as-a-service distributions.

SANDVINE'S CLOUD CAPABILITIES

Sandvine's cloud-based design allows operators to deploy the rich and diverse use cases within the ANI portfolio. This support positions 5G operators to reap the performance, scale, and agility benefits of the cloud while assuring and optimizing their networks.

Two major differentiating characteristics of the Sandvine portfolio are the industry's:

- **Most accurate traffic classification:** effective and accurate classification—even for encrypted traffic—using advanced machine learning provides a trustworthy foundation for our full range of use cases
- **Largest collection of valuable out-of-the-box use cases:** Sandvine's customers reap the rewards of dozens of proven Analytics, Network Optimization, Revenue Generation, Revenue Assurance, Regulatory Compliance, and Automation use cases

Sandvine's cloud-based solution consists of a collection of orchestration templates, VNF, Network Services Descriptors, and VNF images developed for automated deployment and lifecycle management in cloud environments.

Figure 1

Sandvine's simplified and flexible architecture enables cost-effective active network intelligence that seamlessly integrates into cloud environments.

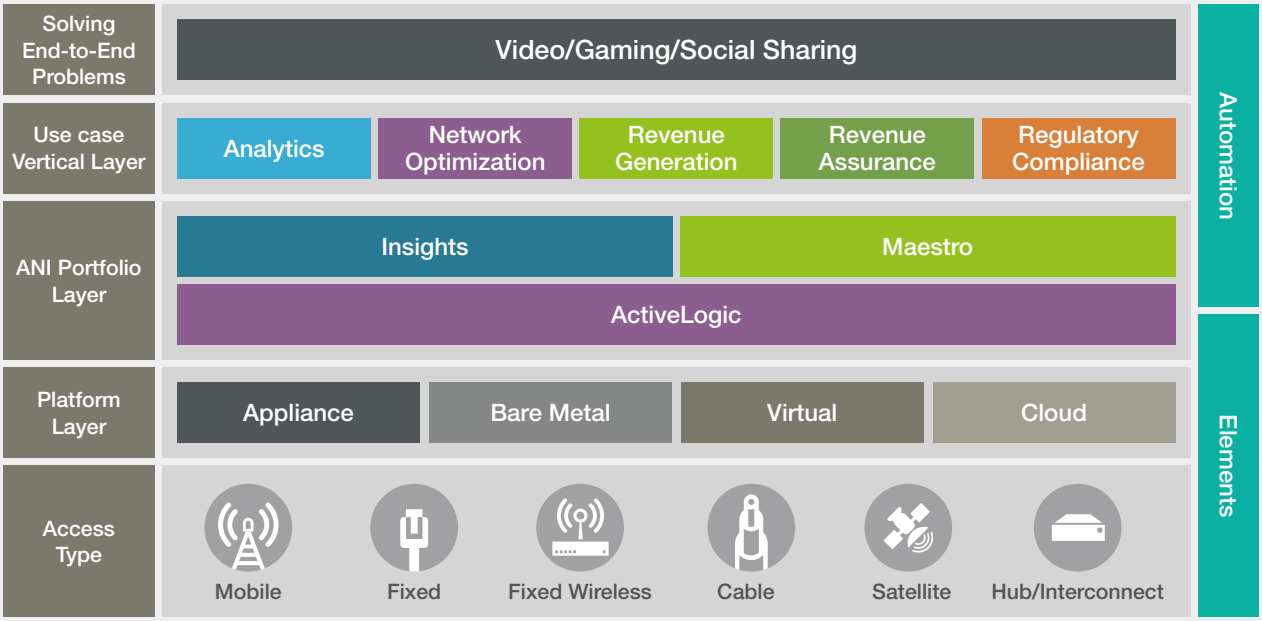

Virtual Network Functions for OpenStack Environments and Hyperscale Cloud Infrastructure

To ease cloud adoption, ETSI has defined a common cloud and Virtual Network Function (VNF) standard designed to facilitate operators in building and running reliable and flexible multi-vendor networks.

Adhering to these standards and meeting the performance and latency goals with service-oriented solutions is a critical success factor for network deployments in the cloud.

[1] For more information, please refer to our Advanced Traffic Classification Whitepaper

Figure 2

Sandvine's cloud architecture is compatible with deployments that comply with or resemble the NFV architecture defined by ETSI

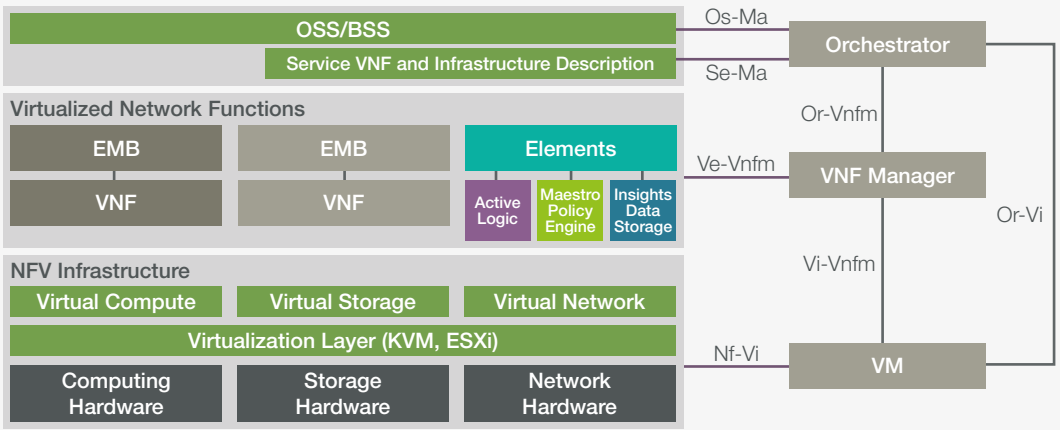

Sandvine VNFs support various lifecycle operations, including VNF/NS (Network Service) instantiation, termination, and automation of zero-day configuration; exact lifecycle scenarios depend on NFV Orchestrator capabilities

Each component of Sandvine's ANI Portfolio is represented by VNFs, consisting of a VNF Descriptor and VNF Image, and includes these products in VNF format:

- **ActiveLogic VNF:** Hyperscale performance data plane, machine learning-powered advanced traffic classification—ready for any access or network
- **Maestro Policy Engine VNF:** Contextually aware control plane with flexible policy and charging control enabling application-based charging
- **Insights Data Storage VNF:** Analytics backend, high-performance database
- **Deep Insights VNF:** Analytics frontend, ANI Portal, and use case visualizations

Elements VNF: Operation and maintenance of ANI Portfolio

OpenStack Support

Within a virtualized network architecture, orchestration coordinates the resources and networks needed to set up cloud-based services and applications.

Sandvine releases cloud orchestration packages for every validated NFV orchestration solution. Cloud orchestration packages consist of VNF and NS Descriptors and meta-information needed for deployment of the Sandvine so lution.

Sandvine VNFs can be deployed in the latest OpenStack environments supplied by major NFVI/VIM vendors.

Table 1
Sandvine ANI VNFs are vertically scalable across all deployment sizes, small to large, allowing operators to optimize footprint based upon deployment needs. This table summarizes performance test results for a representative large deployment in private cloud OpenStack.

	ActiveLogic VNF	Maestro Policy Engine VNF
vCPUs	46 vCPU	16 vCPU
RAM	178 Gb	64 Gb
Storage	80 Gb	200 Gb
Hypervisor Type and Version	QEMU/KVM version 2.11	
Hardware type	Dell R740 with Intel® Xeon® Platinum CPU	
OpenStack EPA Settings	HugePages, CPU pinning andSRIOV interfaces	
Total Throughput	Up to 90 Gbps	
Concurrent Active Connections	Up to 2,500,000 conccurrent active connections	
Provisioned Subscribers	Up to 20,000,000 total subscribers	

Rather than force operators to build their own network functions virtualization (NFV) solutions from scratch, many companies are teaming up to offer pre-integrated solutions. One prominent example is the Cisco NFV Infrastructure Solution. This ETSI-compliant solution uses the Red Hat OpenStack platform and Intel hardware to deliver carrier-grade high availability, reliability, and predictable performance.

Sandvine's cloud architecture has achieved functional, performance, and orchestration validation on the Cisco NFV Infrastructure Solution, including on Cisco:

- VIM (Virtualized Infrastructure Manager)
- UCS (Unified Computing System)
- ESC (Elastic Services Controller)
- NSO (Network Services Orchestrator)

Figure 3

Sandvine's cloud architecture is validated on the Cisco NFV Infrastructure Solution, which incorporates the Red Hat OpenStack platform and Intel hardware

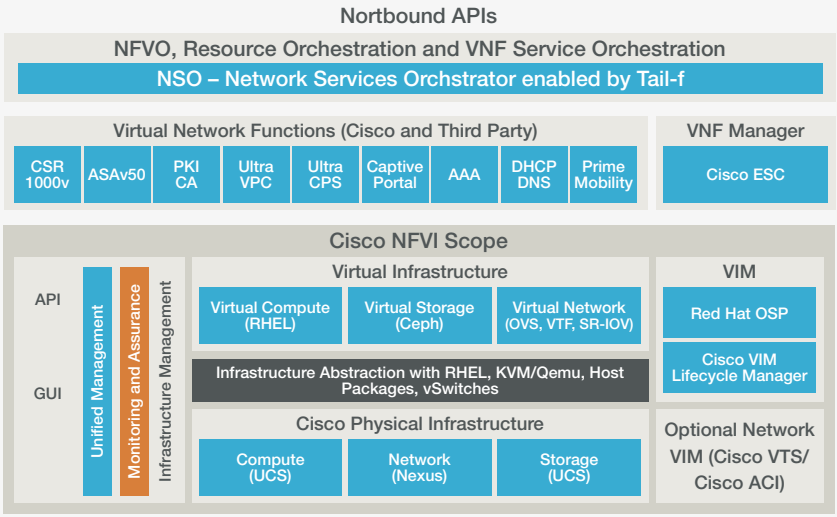

Hyperscale Cloud Compatibility

Public cloud deployments on hyperscale Infrastructure-as-a-Service (e.g., AWS, GCP, Azure) is an essential part of many operators' plans for next-generation networks, particularly with the build-out of 5G.

Sandvine's architecture is completely compatible with hyperscale cloud deployments. Moreover, to simplify installation and management in cloud environments, the Sandvine solution:

- Supports single-click deployment of cloud services
- Uses common orchestration tools (e.g., Terraform) to support cross-platform deployments and eliminate platform dependencies
- Provides robust user, network, and instance security

Example: Amazon Web Services

Sandvine's cloud architecture supports a range of deployment options (see Figures 4-5) through deployment on the Amazon public cloud.

These deployment scenarios present a range of options to deploy Sandvine's data plane (ActiveLogic), control plane (Maestro Policy Engine), datastore (Insights VNF), ANI Portal (Deep Insights VNF), and operations, administration, and maintenance (OAM, the Elements VNF) to enable Application and Network Intelligence in networks that use cloud technologies.

Figure 4

In this configuration, the Sandvine ANI Portal, datastore, and OAM functions are deployed in the public cloud on Amazon EC2 instances; Sandvine control plane and data plane functions are deployed in the operator's datacenter on proprietary hardware

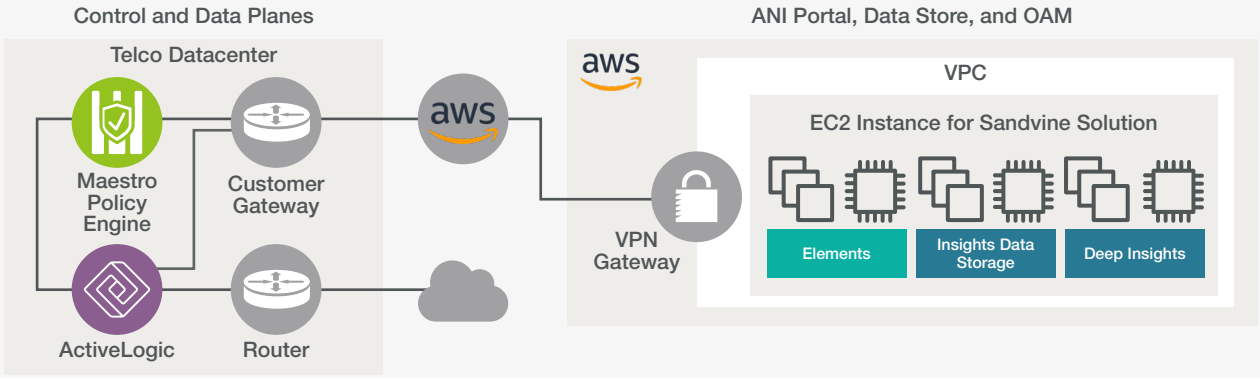

Figure 5

In the full public cloud deployment, all Sandvine elements and functions—ANI Portal, datastore, OAM, control plane, and data plane—are deployed in the public cloud on Amazon EC2 instances

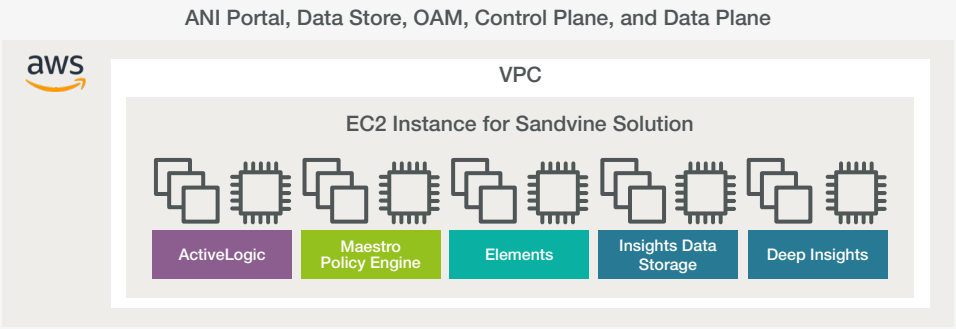

Cloud-native Service Design

Cloud-Native solutions are especially appealing to operators for reasons of cost optimization. Mobile networks are predominantly access-and-edge networks, with the cost function being heavily influenced by the number of base stations and their necessary aggregation.

At the same time, market forces—particularly in the form of increased competition—are placing pressure on average revenue per user (ARPU); accordingly, it is imperative to bring down the costs associated with serving each user.

The hyperscale infrastructure model offers significant promise to operate massive network infrstructure at a very lean OpEx, due primarily to the ability to automate most operations and deployments while maintaning high availability and efficient scale. Additionally, the improved flexibility is expected to enable network operators to bring services to market faster and with lower expenses.

To allow operators to fulfill the promise of 5G and cloud-native networks, Sandvine follows the design principles laid out by Cloud Native Computing Foundation (CNCF) for Cloud-native Network Functions:

- **Compatibility:** Sandvine CNFs work with Certified Kubernetes products and CNI-compatible network plugins
- **Statelessness:** Sandvine CNFs manage state by decoupling application and data storage logic and are resilient to node failures
- **Security:** Sandvine CNFs use security and network policy to isolate containers from one another and the host
- **Scalability:** Sandvine CNFs support horizontal scaling across multiple machines
- **Observability:** Sandvine CNFs externalize their internal states in a way that supports metrics, tracing, and logging
- **Installable and Upgradeable:** Sandvine CNFs use standard, in-band deployment tools such as Helm charts
- **Hardware Resources and Scheduling:** Sandvine CNF containers use device plugins to access all hardware and schedule to specific worker nodes depending on hardware capabilities

ABOUT SANDVINE

Sandvine's cloud-based Application and Network Intelligence portfolio helps customers deliver high quality, optimized experiences to consumers and enterprises. Customers use our solutions to analyze, optimize, and monetize application experiences using contextual machine learning-based insights and real-time actions. Market-leading classification of more than 95% of traffic across mobile and fixed networks by user, application, device, and location creates uniquely rich, real-time data that significantly enhances interactions between users and applications and drives revenues. For more information visit **http://www.sandvine.com** or follow Sandvine on Twitter **@Sandvine**.

SANDVINE

USA
5800 Granite Parkway
Suite 170
Plano, TX 75024
USA

EUROPE
Svärdfiskgatan 4
432 40 Varberg,
Halland
Sweden
T. +46 340.48 38 00

CANADA
410 Albert Street,
Suite 201, Waterloo,
Ontario N2L 3V3,
Canada
T. +1 519.880.2600

ASIA
RMZ Ecoworld,
Building-1, Ground Floor,
East Wing Devarabeesanahalli,
Bellandur, Outer Ring Road,
Bangalore 560103, India
T. +91 80677.43333